CHAPTER

6

ACTIVE DIRECTORY CONFIGURATION

After completing this chapter, you will be able to:

- Create sites within Active Directory
- Create subnets within Active Directory
- Create site links within Active Directory
- Create site link bridges within Active Directory
- Create connection objects within Active Directory
- Create Global Catalog Servers within Active Directory
- Move server objects between sites within Active Directory
- Transfer operations master roles within Active Directory
- Implement an Organizational Unit structure

This chapter discusses the steps necessary to create an Active Directory structure for an enterprise network environment. The Active Directory elements that we will examine in this chapter include sites, subnets, site links, site link bridges, connection objects, Global Catalog Servers, operations masters, and more. We will look at how each element is configured within Windows 2000 and how these elements interact.

Creating a Site

As we discussed in Chapter 2, **sites** are collections of computers that are connected via a high-speed network. Typically, the computers within a site are connected via local-area network (LAN)–style technology and are considered to be well connected. **Well connected** generally means constant high-speed connectivity within an IP subnet, although a site can include multiple subnets.

Windows 2000 creates the first site automatically when Active Directory is installed. This site is named Default-First-Site, and it includes the initial domain controller (DC). For a small LAN, the single site will be sufficient. For larger environments, however, additional sites must be created manually. To create a site, open the AD Sites and Services snap-in, shown in Figure 6-1, and open the context menu of the Sites folder. Select the New Site option to create a new site.

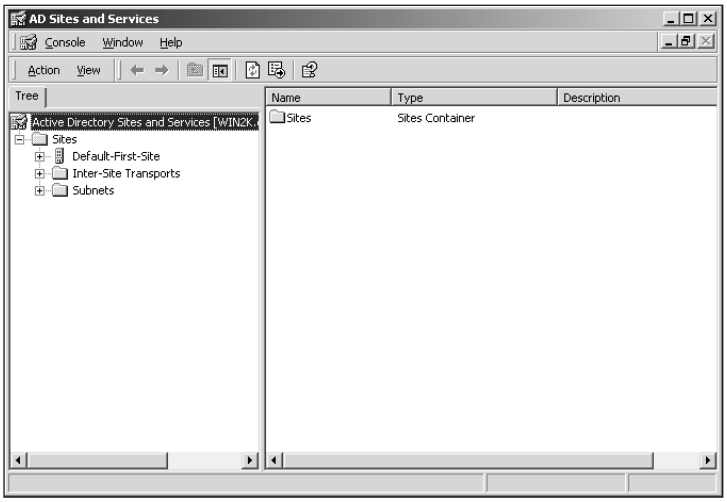

Figure 6-1 AD Sites and Services snap-in

The New Object–Site screen, shown in Figure 6-2, allows you to enter the name of the remote site and to select the site link for the site. Windows 2000 creates a default site link called DEFAULTIPSITELINK that can be used to establish the replication process of the Active Directory service. This default site link uses remote procedure calls (RPCs) over TCP/IP, and will use any available route to the remote site for replication. If explicit site links have been previously defined, those site links will show up in the lower portion of the New Object–Site screen.

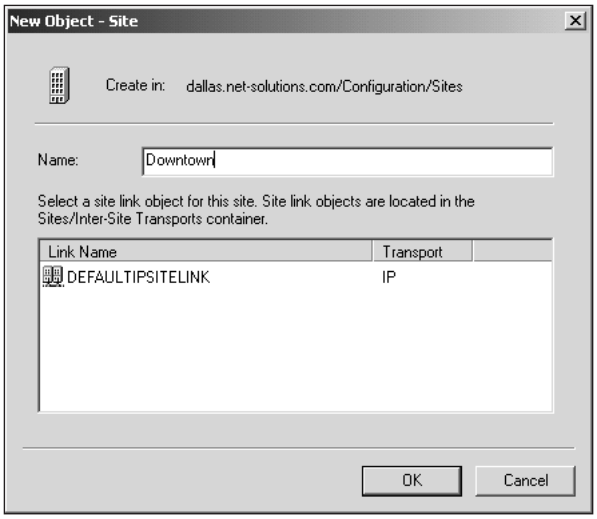

Figure 6-2 Creating a new site

Once the site is defined, you must undertake several other steps before the site can be activated within the Active Directory structure. These steps are nicely delineated in the dialog box that follows the creation of a new site, as shown in Figure 6-3. To finish configuring a site, you must do the following:

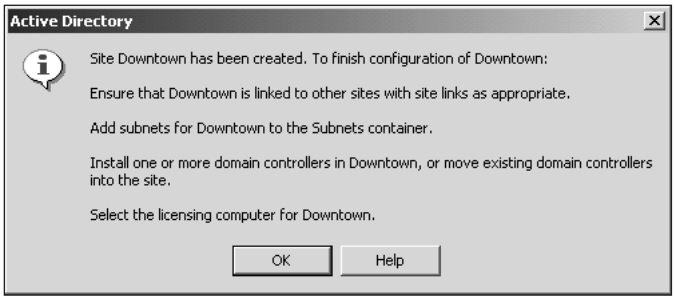

Figure 6-3 Required configuration steps for a new site

- Add appropriate IP subnets to the site.
- Install or move a domain controller or controllers into the site. Although a DC is not required for a site, it is strongly recommended.
- Connect the site to other sites with the appropriate site link.
- Select a server to control and monitor licensing within the site.

After these steps are completed, the site is added to the Active Directory structure, and the replication is automatically configured by Windows 2000. We will discuss each of these steps in this chapter.

Adding Subnets to Active Directory

Defined sites within Windows 2000 allow for a more efficient replication process than previous versions of Windows NT provided. Replication within areas of the network that are connected via high-speed connections is optimized to minimize latency and to minimize the time required to update records within the Active Directory. In contrast, replication across slower wide area network (WAN) links is optimized to reduce the required bandwidth and avoid flooding the link to a remote location, although Active Directory updates may require more time to take effect. In order for this more efficient replication process to function, Windows has to understand the network topology.

Because every network environment is different, Active Directory does not attempt to create sites nor to associate the subnets with the sites. Instead, the network administrator is tasked with creating these sites and associating IP subnets with the sites. We have already discussed creating a site, so we will now look at associating a subnet with these sites.

To associate a site with a subnet or a group of subnets with a site, first select the Subnets folder from the AD Sites and Services snap-in. Choose New Subnet from the context menu, as shown in Figure 6-4.

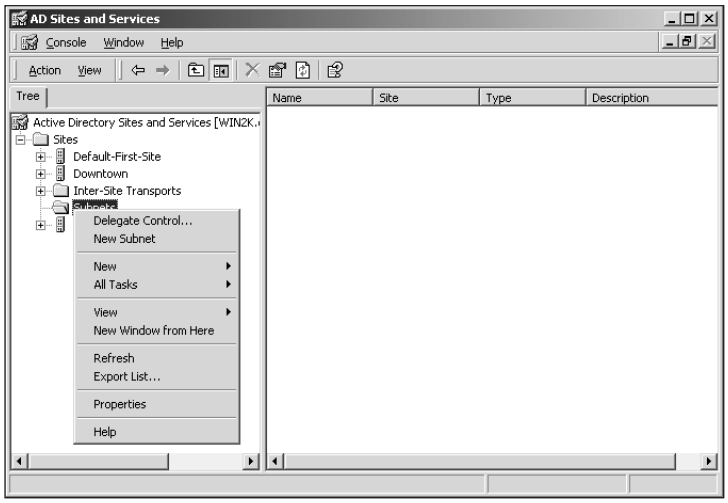

Figure 6-4 Adding a new subnet in the AD Sites and Services snap-in

The new subnet requires the subnet address and the network mask. You enter these items in dotted octet format, which is automatically translated into the network/bit-mask format. You select the site that will be associated with the subnet in the lower section of the property page, shown in Figure 6-5.

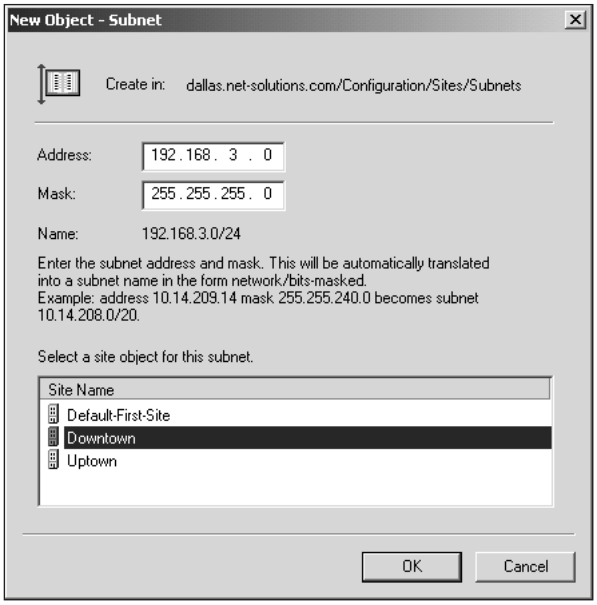

Figure 6-5 New subnet properties

Once the subnet is created and associated with a site, it and other sites will appear in the Subnets folder in the AD Sites and Services snap-in, as shown in Figure 6-6. The properties of each subnet can be manipulated by selecting a subnet and then selecting Properties from the context menu.

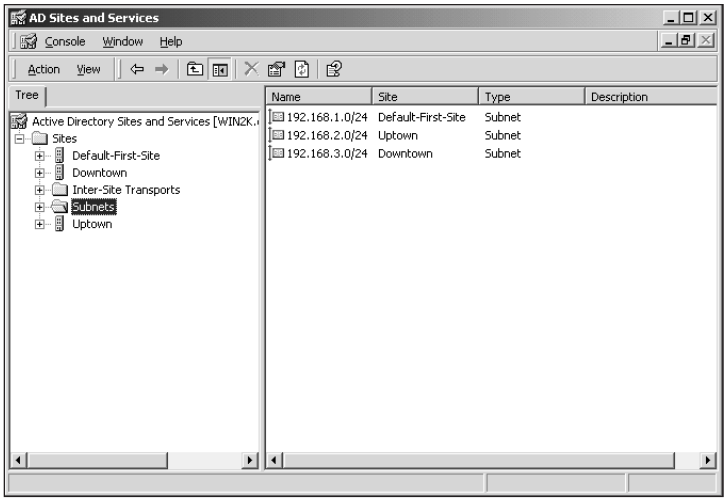

Figure 6-6 Subnets within the AD Sites and Services snap-in

Site Connections

A site is a subnet or selection of subnets connected via a high-speed connection. The sites themselves are connected via site links. **Site links** are low bandwidth or unreliable/occasional connections between sites. In general, any connection between locations slower than LAN speeds is considered a site link. WAN links such as frame relay connections are examples of site links, as are high-speed links that are saturated and have a low effective bandwidth.

Site links are not automatically generated by Windows 2000. Instead, the administrator creates the site links through the AD Sites and Services snap-in. The site links are the core of Active Directory replication. The links can be adjusted for replication availability, bandwidth costs, and replication frequency. Windows 2000 uses this information to generate the replication topology for the sites, including the schedule for replication.

Windows 2000 DCs represent the inbound replication through a special object known as a **connection object**. Active Directory uses site links as indicators for where it should create connection objects, and connection objects use the physical network connections to replicate directory information. Each DC creates its own connection objects for replication within a site (**intrasite replication**). For replication between sites (**intersite replication**), one DC within each site is responsible for evaluating the replication topology. The DC creates the connection objects appropriate to that topology. The server that is responsible for evaluating and creating the topology for the intersite replication is known as the Inter-Site Topology Generator (ISTG).

Site links, like trusts, are transitive. This means that DCs in one site can replicate with DCs in any other site within the enterprise through these transitive links. In addition, explicit links can be created to enable specific replication paths between sites.

Creating a Site Link

Windows 2000 creates a default site link named, naturally enough, DEFAULTIPSITELINK. This site link can be used to connect sites in simple network environments, but in more complicated enterprise environments, you should establish explicit site links.

To create a site link, first open the AD Sites and Services snap-in. Open the Inter-Site Transports folder and then right-click on the appropriate transport protocol, as shown in Figure 6-7. Select New Site Link from the context menu to form a new link.

In our example, the name of the new site link is Uptown Frame Link. Although the name of the link is arbitrary, good administrative practice dictates that the name should be something that identifies the link, the connected sites, and the type of link. Of course, the link could be named Bob, but that name would probably confuse successors and co-workers.

The next step is to select the linked sites from the left column in the New Object–Site Link dialog box and click on the Add button to associate them with the link, as shown in Figure 6-8. A link must contain at least two sites; in general, a link will connect only two sites. If multiple sites exist at one physical location or are connected via a particular network path, however, then those sites could share a single site link.

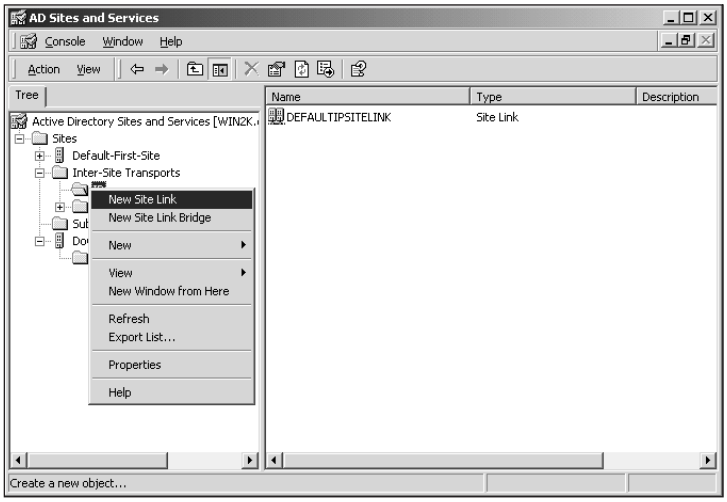

Figure 6-7 Creating a new site link

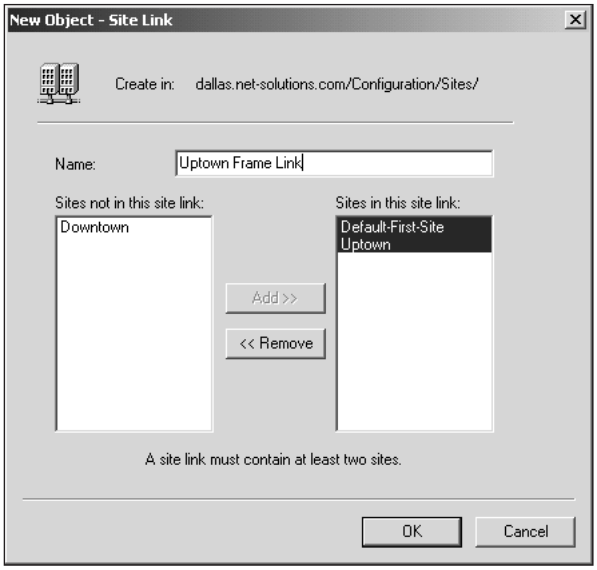

Figure 6-8 Naming the site link and associating the sites

Each site link has four properties that are important, as well as an optional descriptor. The properties are as follows:

- *Name:* A name that uniquely identifies the site link. As discussed earlier, this name should clearly indicate the sites being linked and the speed/type of circuit.

- *Cost:* The relative speed of the link in relation to the other links within the topology. The cost has nothing to do with the actual monetary cost of the bandwidth. Links with lower costs are faster, whereas links with higher costs are slower. The cost defaults to 100 on a new circuit.
- *Transport:* Indicates the type of transport used to replicate the directory information between the DCs. You have two options: synchronous RPC over a routed TCP/IP connection, or an asynchronous Simple Mail Transfer Protocol (SMTP) connection over the underlying mail transport network. This property is not set within the link properties, but is instead determined when the site link is first created.
- *Schedule:* Determines when the directory information is replicated between sites. This property is determined by two elements: the replication frequency and the available times. The replication frequency is adjusted within the properties of the site link, as shown in Figure 6-9. The schedule is a list of times that the site link is available to pass replication data. It is adjusted through the Change Schedule option within the site link properties.

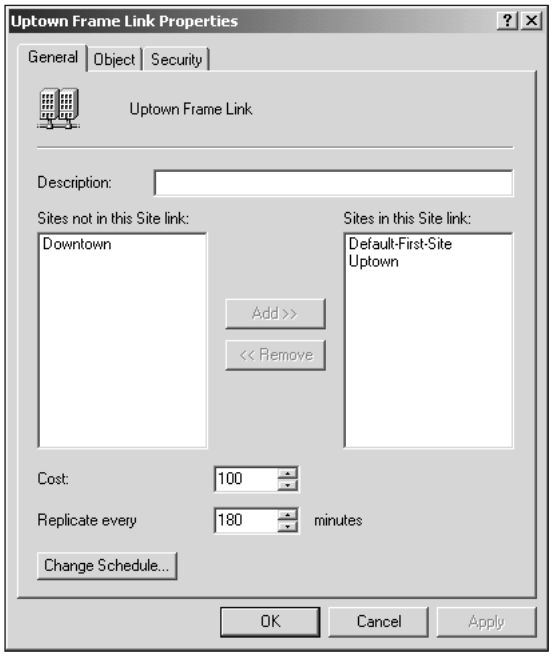

Figure 6-9 Site link properties

Site Link Bridges

Within a fully routed network, site links are transitive. As a result, all the site links for a particular transport are bridged together, and the replication can route between sites as needed. By default, Windows 2000 bridges all the site links for a particular transport.

If the network is not fully routed, then site link bridges must be explicitly defined for each transport. The transitive link feature can be turned off within each transport. You do so by unselecting the Bridge All Site Links option within the property sheet of each transport, as seen in Figure 6-10. Once the option is unselected, site link bridges allow transitive replication routing within the bridged links, but not outside the bridge.

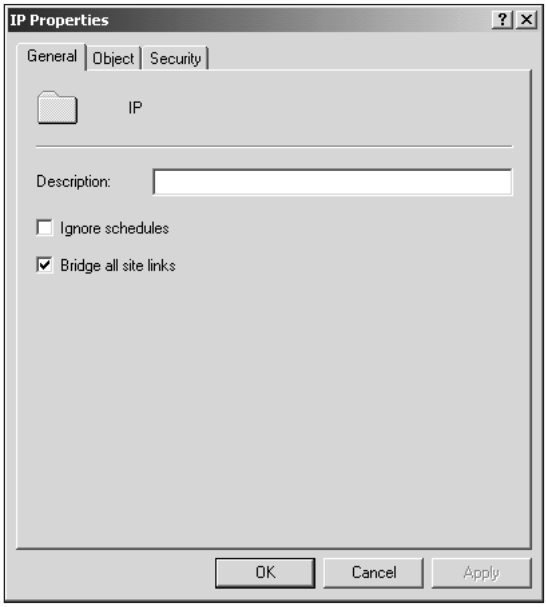

Figure 6-10 Default site link bridging

To create a new site link bridge, first you must have defined the site links themselves, as discussed earlier. Then, open the Inter-Site Transports folder and select the desired transport. This transport can be either IP or SMTP. From the context menu of the selected transport, choose New Site Link Bridge, as shown in Figure 6-11.

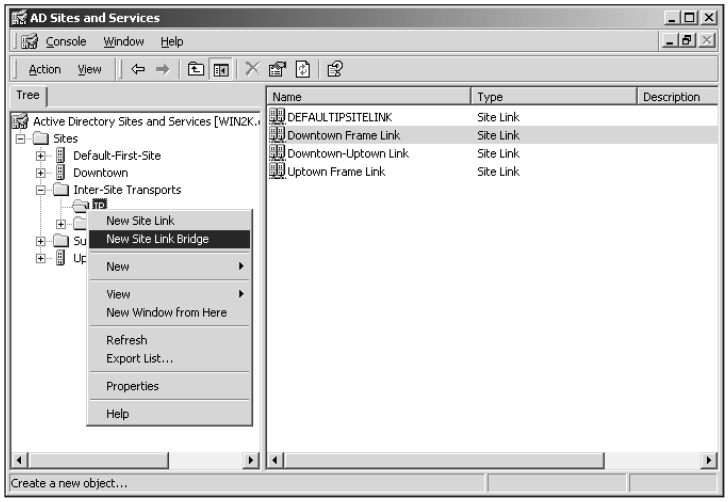

Figure 6-11 Creating a new site link bridge

The new site link bridge requires at least two site links. When these site links are bridged, a transitive replication link is generated across both links. In the case of our example in Figure 6-12, the downtown link connects the corporate center with the downtown center, and the uptown link connects the corporate center with the uptown center. Through the site link bridge, the uptown site can now replicate directly with the downtown site, even though no direct physical link exists between the sites.

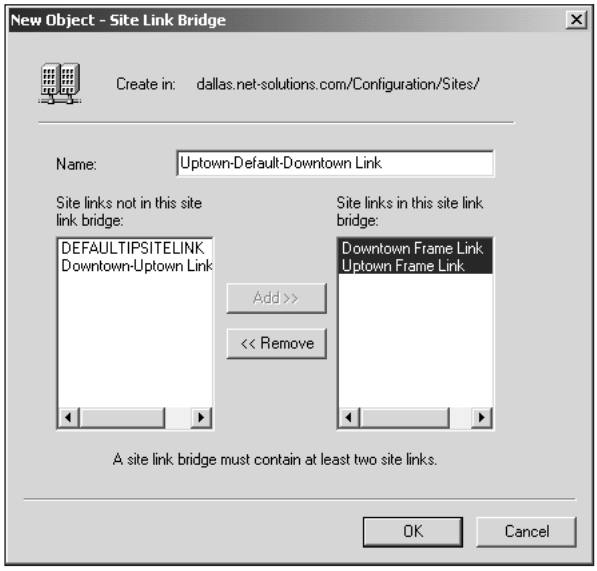

Figure 6-12 The downtown link connects the corporate center with the downtown center, and the uptown link connects the corporate center with the uptown center

Connection Objects

As we explained earlier, Windows 2000 DCs represent inbound replication through a special object known as a **connection object**. In general, connection objects will be automatically generated both within sites and between sites. If you don't use the Knowledge Consistency Checker (KCC) to generate replication topology, however, then you will have to generate the connection objects manually. To generate these objects, first open the AD Sites and Services snap-in, and then navigate to the DC on which you wish to form a connection object. Select NTDS Settings, and then choose New Active Directory Connection from the context menu, as shown in Figure 6-13. The next screen, shown in Figure 6-14, lists the DCs that can be used for inbound replication.

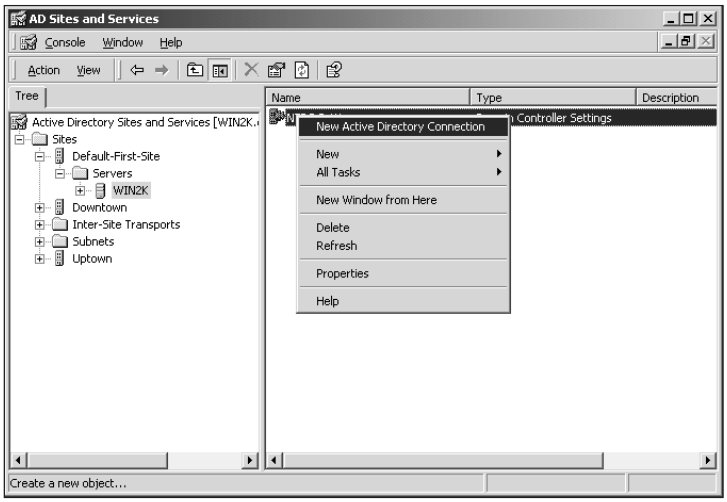

Figure 6-13 Creating a new connection object

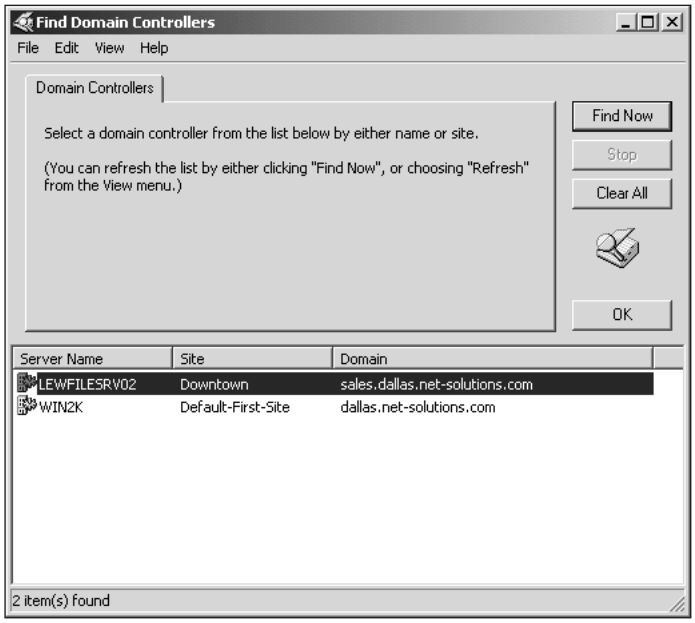

Figure 6-14 Available DCs for replication

Choose the desired DC and then click on OK. The selected DC appears in the next screen, as shown in Figure 6-15. If this is the correct DC for inbound replication, click on OK to finish creating the site connector.

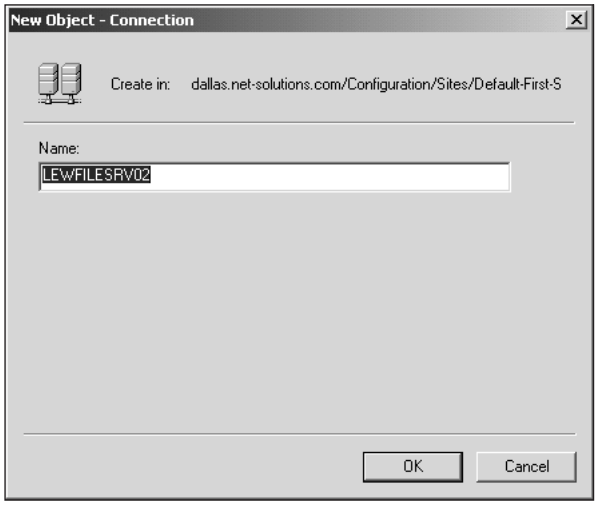

Figure 6-15 Finish creating the connection object

Once the connection object is created, you can view it through the NTDS Settings for each DC, as seen in Figure 6-16. Remember that connection objects are usually created automatically, and they can be dynamically modified to change replication if new DCs and sites are created. A connection object should be manually created only if you are absolutely confident that it will be needed on a permanent basis.

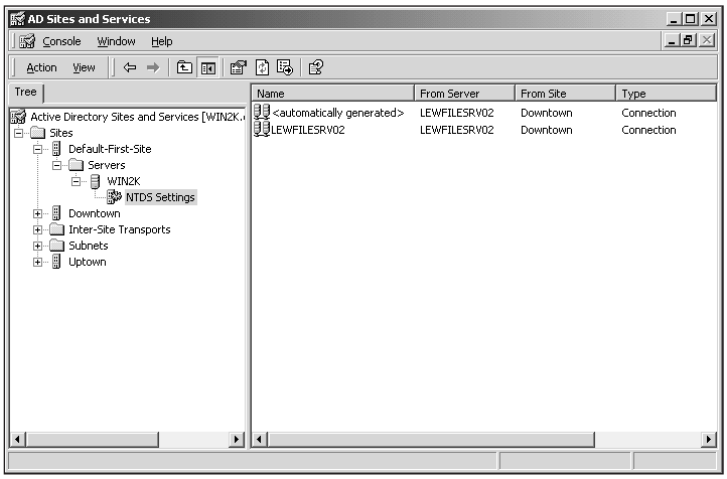

Figure 6-16 Viewing connection objects

Manually created connection objects will remain until manually deleted.

Both manually and automatically created connection objects can be viewed through the NTDS Settings for each of the DCs. In the previous example, you will notice that the manually created object and the automatic object are identical, because we had only two DCs to work with.

Moving Domain Controllers between Sites

We discussed earlier how subnets can be associated with particular sites. After a site has been associated with a subnet, any new DC with an IP address within that subnet will automatically be assigned to the site. For example, if site Downtown has the subnet 192.168.1.0/24 associated with it, a new DC with the IP address 192.168.1.5 will automatically become part of the Downtown site. If a DC is assigned an IP address that is not associated with a particular site, the new DC will be assigned to the default site.

In some situations, the automated assignment does not fit the needs of the network environment, or pre-existing DCs need to be moved to the correct sites. Fortunately, this is a very easy process.

To move a DC between two sites, first open the AD Sites and Services snap-in. Navigate to the server that you wish to relocate, and then open the context menu for that server, as shown in Figure 6-17.

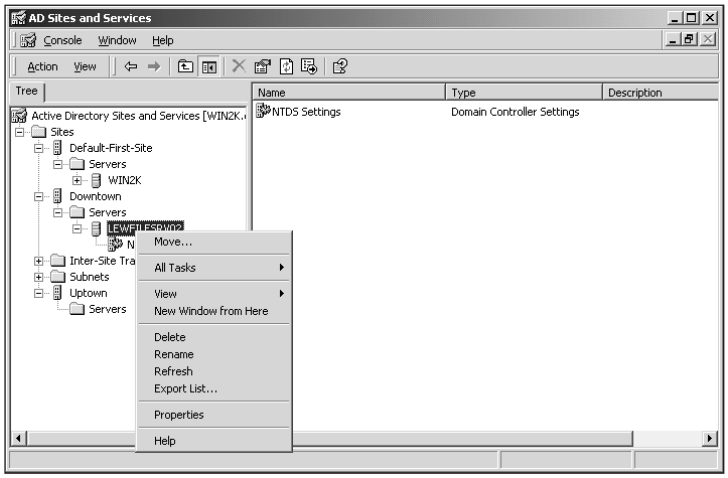

Figure 6-17 Moving a DC between sites

Select Move from the context menu, and then select the destination site, as shown in Figure 6-18.

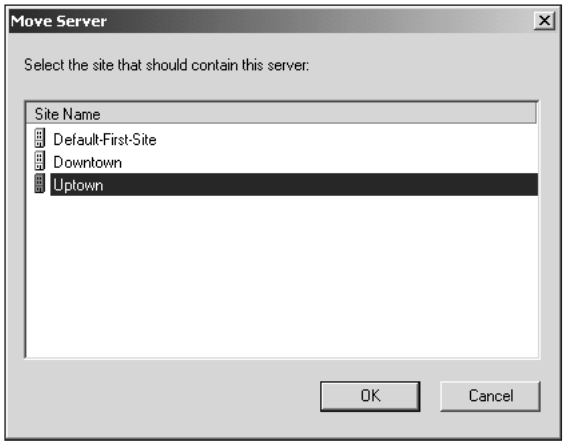

Figure 6-18 Select the destination site for the DC

Click on OK to move the server to the destination site. Obviously, doing so does not change the actual network settings on the DC itself. If IP address changes or other network configuration changes are necessary, those changes will need to be made on the DC before it will be able to communicate with the rest of the network environment.

The new configuration of the site will automatically be displayed within the AD Sites and Services snap-in, as shown in Figure 6-19. Note the new location of the moved DC.

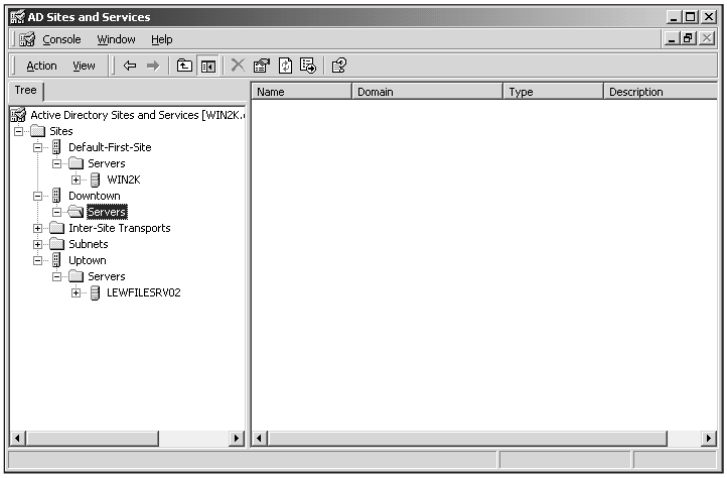

Figure 6-19 New site location for the DC

Global Catalog Servers

DCs keep a complete copy of the Active Directory database for that domain, so that information about each object in the domain is readily available to the users and services. This setup works well for the local domain, but what about information for other domains and the objects within those other domains? Remember, one of the design goals for Windows 2000 was a unified logon, no matter where a user was located within the domain tree. Obviously, for such a unified logon to work, the local DCs must have some information about the other domains within the tree and forest. However, replication of all the information about all the objects in all the domains within a forest simply isn't feasible.

Windows 2000 solves this issue through the use of a special limited database. This database is known as the Global Catalog. The Global Catalog stores partial replicas of the directories of other domains. The catalog is stored on DCs that have been designated as Global Catalog Servers. These servers also maintain the normal database for their domain.

Function of the Global Catalog

The Global Catalog has two primary functions within Active Directory. These functions relate to the logon capability and the queries within Active Directory. We will examine each in detail next.

Within a native-mode multidomain environment, the Global Catalog is required for logging in to the network. The Global Catalog provides universal group membership

information for the account that is attempting to log on to the network. If the Global Catalog is not available during the logon attempt and the user account is external to the local domain, the user will be allowed to log in to only the local machine.

Obviously, if the account is part of the local domain, the DCs for the local domain will handle the authentication request. The Global Catalog is required only when a user account or object needs to be authenticated by another domain.

The majority of Active Directory traffic consists of queries, and queries for objects (printers, services, and so on) occur much more often than database updates. Within a simple single-domain environment, the directory is readily available for these queries. However, imagine for a moment a highly complex multidomain environment. It doesn't make any sense to require every query to search through each domain.

The Global Catalog maintains a subset of the directory information available within every domain in the forest. This process allows queries to be handled by the nearest Global Catalog, and thus saves time and bandwidth. If more than one DC is a Global Catalog Server, the response time for the queries improves. Unfortunately, each additional Global Catalog Server increases the amount of replication overhead within the network.

The Global Catalog is a read-only database, unlike the normal directory database.

Creating Global Catalog Servers

Windows 2000 automatically creates a Global Catalog on the first DC within a forest. Each forest requires at least one Global Catalog. In an environment with multiple sites, it is good practice to designate a DC in each site to function as a Global Catalog Server. Remember, native-mode Windows 2000 domains require a Global Catalog to allow users to complete the authentication process and log in to the network. A mixed-mode domain does not require a Global Catalog Server.

If additional Global Catalog Servers are desired, the service can be added to any DC within the forest. You do so through the AD Sites and Services snap-in. The Global Catalog can also be moved off the initial DC if additional DCs are available.

To create a Global Catalog Server, first open the AD Sites and Services snap-in. Then, navigate to the DC that you wish to be a Global Catalog Server. Highlight NTDS Settings for the desired server and then select Properties from the context menu, as shown in Figure 6-20. Doing so will bring up the screen shown in Figure 6-21.

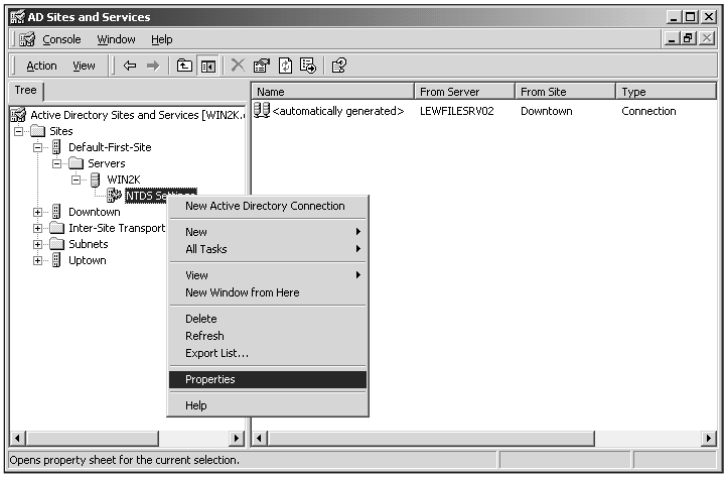

Figure 6-20 Creating a Global Catalog

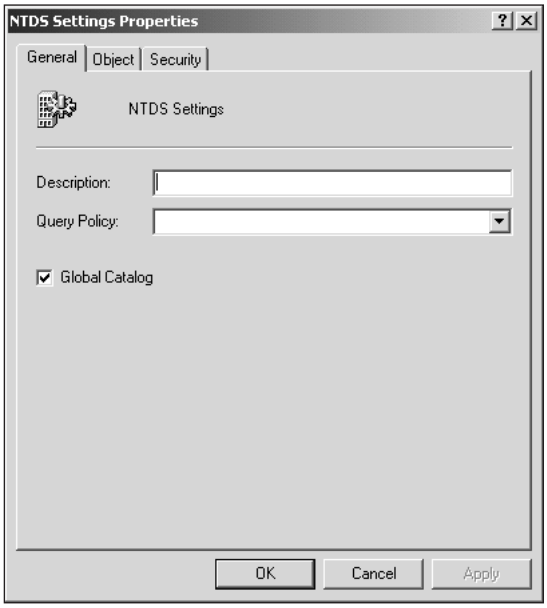

Figure 6-21 NTDS Settings

To enable the DC as a Global Catalog Server, simply check the box labeled Global Catalog. Deselecting the box will remove the Global Catalog from the DC.

Although any and all DCs can be configured as Global Catalog Servers, a sense of balance is necessary when designating these servers. As the number of Global Catalog

Servers increases, the response time to user inquiries decreases. However, the replication requirements within the environment increase as the number of Global Catalog Servers increases.

Operations Masters

Much of the replication within an Active Directory environment is multi-master replication, which means that the DCs are all peers. This is in contrast to earlier versions of Windows NT, in which a Master Domain Controller was responsible for recording all changes to the security policy and replicating those changes to the backup DCs.

Several types of operations are impractical for a multi-master environment, however. Windows 2000 handles these operations by allowing only a single DC to make these types of changes. This DC is known as an **operations master**. Five different operations master roles can be assigned to DCs: the schema master, domain naming master, relative ID master, Primary Domain Controller (PDC) emulator, and infrastructure master. Each of these roles will be discussed in detail a bit later.

The schema master and domain naming master operations master roles are assigned on a forestwide basis. Only one of each operations master can exist within a forest.

Schema Master

The schema master controls all the updates and modifications to the schema itself. As you will recall, the schema controls the definition of each object in the directory and its associated attributes.

To change the schema master, begin by opening the Active Directory Schema Manager. If this tool has not been installed on the DC, install the Schema Manager from the Windows 2000 CD-ROM.

Select the Active Directory Schema Manager. Right-click on it and choose Change Domain Controller, as shown in Figure 6-22, to select the desired schema master. You will have two options: The first allows Active Directory to select the new DC/schema master, and the second allows you to select the target DC. For our purposes, we will designate the new schema master. Enter the name of the target DC, as shown in Figure 6-23.

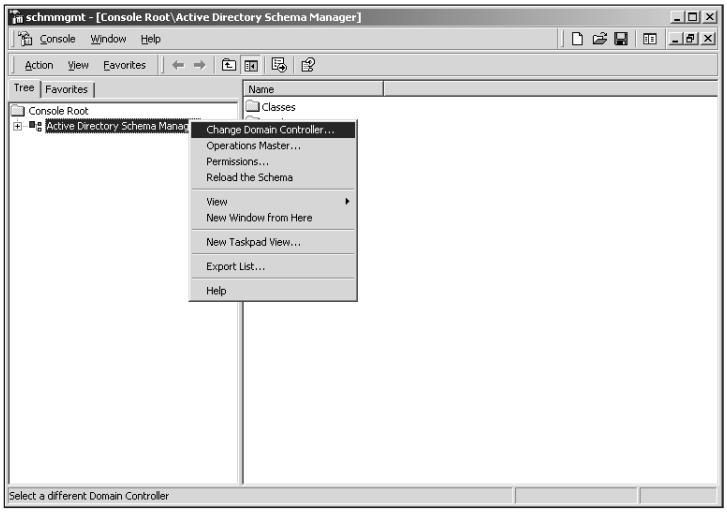

Figure 6-22 Changing the schema master

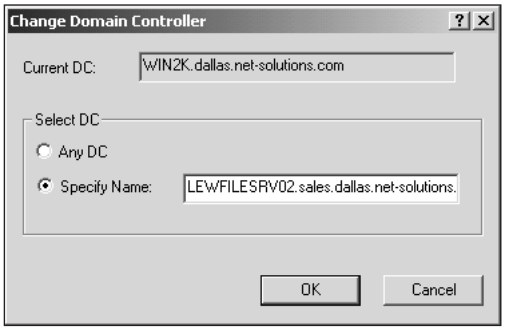

Figure 6-23 Selecting the desired DC

Doing so will change the focus of the Active Directory Schema Manager to the newly chosen DC. To complete the change, select the Active Directory Schema Manager from the left pane and then choose Operations Master from the context menu, as seen in Figure 6-24.

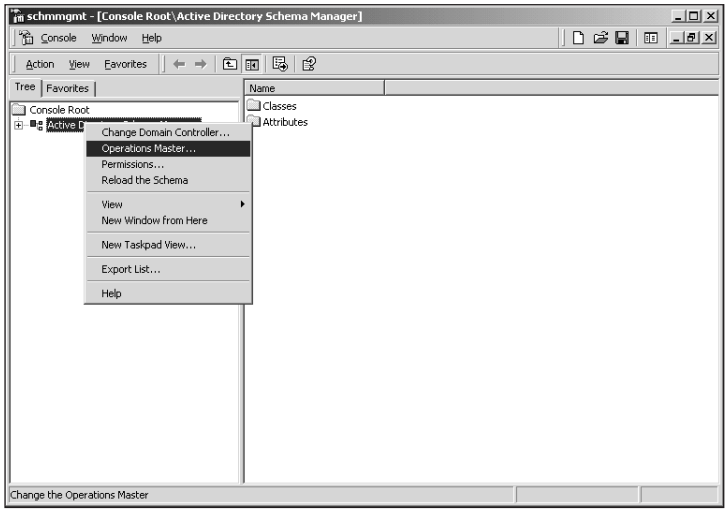

Figure 6-24 Changing the schema master for a forest

Doing so will bring up the property sheet shown in Figure 6-25, which shows the current focus of the Active Directory Schema Manager and the current operations master for the forest. Click on Change to move the schema master to the server listed at the top of the property page from the server listed at the bottom of the property page.

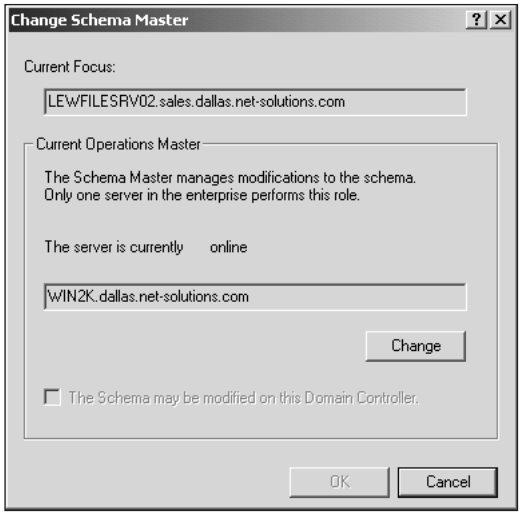

Figure 6-25 Final selection for the schema master

Domain Naming Master

The domain naming master controls the addition of domains to or removal of domains from the forest. As with the schema master, only one domain naming master can exist within a forest.

To modify the domain naming master, begin by opening the Active Directory Domains and Trusts snap-in. This modification is much like the schema master modification. First, connect to the DC to which you wish to transfer the operations master. You do so by selecting Connect To Domain Controller from the context menu, as shown in Figure 6-26.

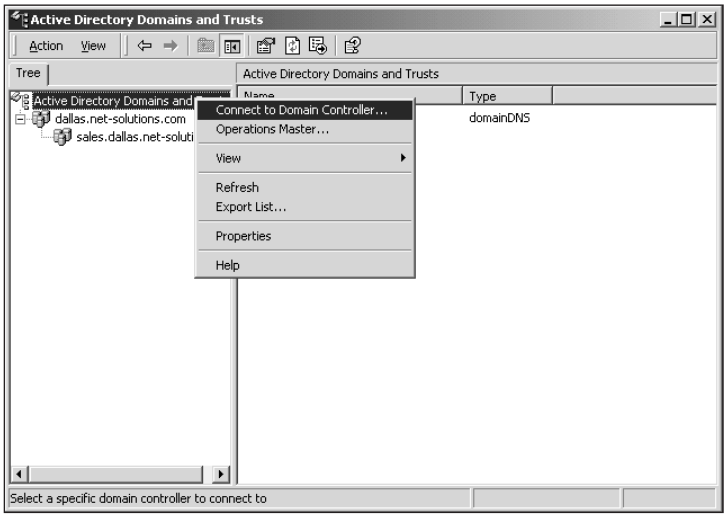

Figure 6-26 Changing the focus of the Active Directory Domains and Trusts snap-in

Select the target domain via the Browse button, and then choose an available DC from the list at the bottom of the screen. Selecting Any Writable Domain Controller will allow Active Directory to choose the target DC. If you wish, you can select a particular DC to which to connect. In Figure 6-27, we chose the DC for sales.dallas.net-solutions.com.

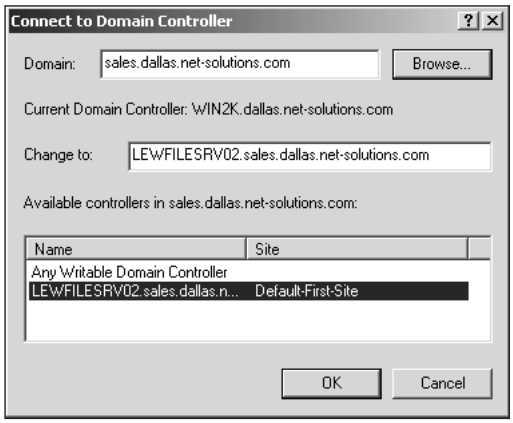

Figure 6-27 Selecting the target DC

Selecting the new DC automatically changes the focus of the Active Directory Domains and Trusts snap-in to the new DC. To change the domain naming master for the forest, select Operations Master from the context menu within the snap-in, as shown in Figure 6-28.

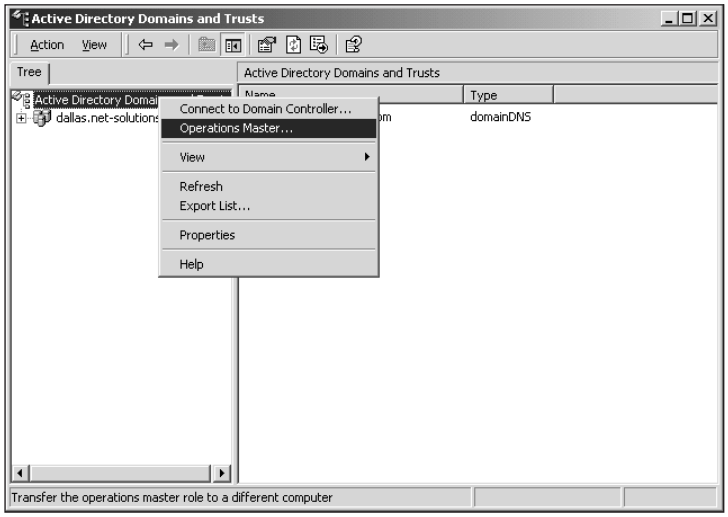

Figure 6-28 Changing the domain naming master for a forest

This selection will bring up a sheet that shows the current domain naming operations master and the targeted DC. Clicking on Change, as shown in Figure 6-29, will move the domain naming master operations role to the DC shown within the second box on the screen. The change takes place instantly, although replication of that change to the other DCs will naturally depend upon the replication topology and schedule.

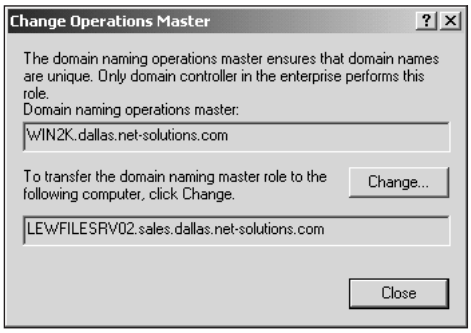

Figure 6-29 Final selection for the domain naming master

Relative ID Master

The relative ID (RID) master controls the sequence number for the DCs within the domain. The master assigns a unique sequence of RIDs to each of the DCs. When the DC uses all the RIDs that the RID master has assigned, the DC receives another sequence of RIDs from the RID master.

By default, the RID master is the first DC installed within a domain. To modify the RID master, first open the Active Directory Users and Computers snap-in. Select Connect To Domain Controller from the context menu, and then navigate to the target DC.

Once the target DC becomes the focus of the Active Directory Users and Computers snap-in, select Operation Masters from the context menu, as shown in Figure 6-30. Doing so will bring up the property sheet shown in Figure 6-31.

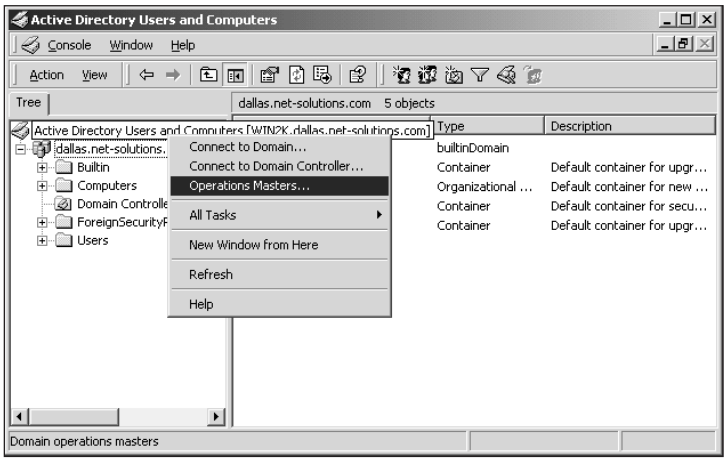

Figure 6-30 Changing operation masters within a domain

Figure 6-31 The RID master settings

The current operations master for the RID service should appear in the upper field, and the target DC will appear in the bottom field. In the case of a domain with only a single DC (not recommended!), the operations master cannot be changed. Click on Change to move the selected role to the new operations master.

PDC Emulator Master

The PDC emulator is used whenever a domain contains non–Windows 2000 computers and acts as a Windows NT PDC for downlevel clients and for Windows NT backup DCs. By default, the first DC installed in a Windows 2000 domain will be the PDC emulator, but this role can be modified if desired.

To change the PDC emulator master, first launch the Active Directory Users and Computers snap-in and change focus to the target DC as described earlier. Then, open the Operations Master properties from the context menu, as shown in Figure 6-30.

Click on the PDC tab to bring the PDC emulator to the forefront of the properties screen, as seen in Figure 6-32. The current operations master will be listed in the upper field, and the target DC will be listed in the lower field. Click on Change to complete the process.

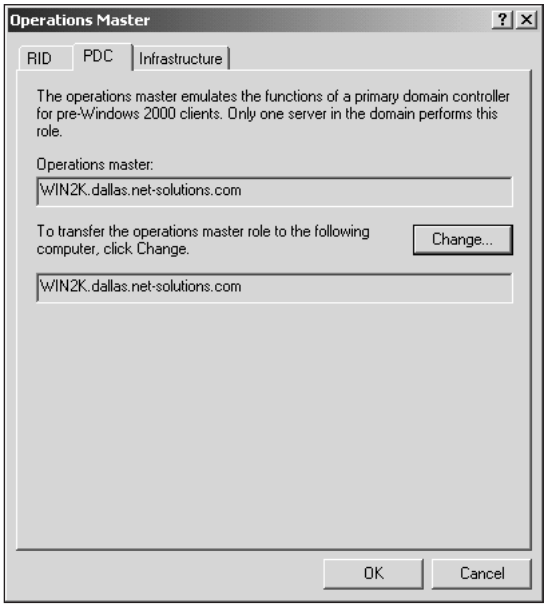

Figure 6-32 The PDC emulator settings

Infrastructure Master

The infrastructure master is responsible for maintaining all interdomain object references. That is, the infrastructure master informs certain objects (such as groups) that other objects (such as users in another domain) have been moved, changed, or otherwise modified. As with the other domainwide operations masters, this role is initially performed by the first DC within a domain. The infrastructure master role can be moved to another DC via the Active Directory Users and Computers snap-in.

To do so, change the focus to the target DC as discussed earlier, and open the operations master property sheet. Click on the Infrastructure tab, as shown in Figure 6-33.

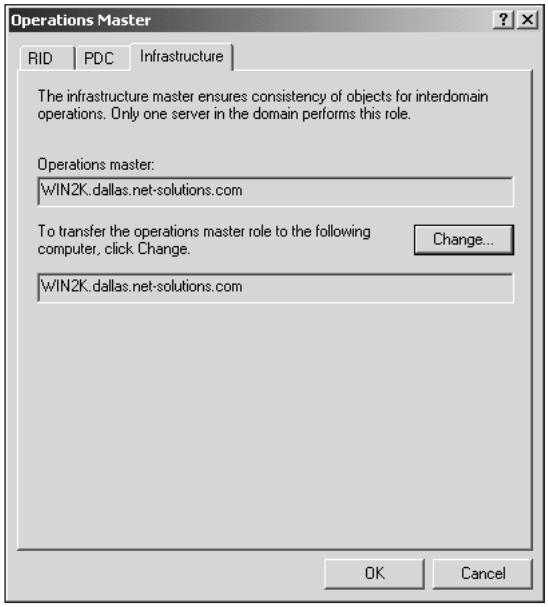

Figure 6-33 The infrastrusture master settings

The current operations master will be listed in the upper field, and the target DC will be listed in the lower field. Click on Change to complete the process.

Implementing Organizational Units

One of the primary advantages of Windows 2000 and the Active Directory service over Windows NT is the ability to control administrative powers more discretely. Under Windows NT, the base unit of administrative power is the domain. There is no way to grant someone administrative power over a subsection of the domain, such as a sales division or geographical office. This limitation means that either you are forced to make every required change to user access rights, or administrative power is granted to a larger circle of people.

Some workarounds exist to this problem, including the user of master domain/resource domain structures; but even these require careful planning and additional infrastructure to function correctly. Particularly annoying is the fact that competing network operating systems offer the ability to segregate administrative roles to a particular element of the network.

Fortunately, Active Directory introduces the Organizational Unit (OU) to the Windows networking environment. An OU is essentially a subset of a domain that can contain any Active Directory object. The network administrator can designate control of and access to each OU and the objects it contains. In addition, policies can be designated on the OUs in order to manage user policies and rights.

Essentially, two main uses exist (so far) for OUs. They are:

- Allow subadministrators control over a selection of users, computers, or other objects.
- Control desktop systems through the use of Group Policy Objects associated with an OU.

We will look at each of these uses in the following sections.

Delegating Control of Part of a Domain

One of the most common administrative needs is the ability to allow others to manage user accounts. A fine line always exists between maintaining security and delegating power to others. Windows NT offers the ability to grant the right to change passwords and other limited administrative control, but these rights apply on a domainwide basis.

Windows 2000 offers a capability to delegate various levels of control on only parts of a domain. You do so using OUs. As discussed earlier, an OU is a container that can contain various objects, including user accounts, computers, printers, shares, services, and much more. An OU can be treated as a subdomain for practical purposes—administrators of a domain retain control of the OU, but specific rights can also be granted to other users or groups.

Follow along as we create an OU and delegate control of it. In our scenario, the marketing department is almost like a separate organization, and its employees have decided that they need the right to change passwords for the division. Tired of changing passwords for marketing people at 2:00 A.M., the IT department agrees.

First, an OU must be created to contain the user accounts and other objects for the marketing department. All OU implementation and administration is accomplished through the Active Directory Users and Computers snap-in. Once the console is started, navigate to the domain within which the OU should be located. From the context menu, choose New | Organizational Unit, as shown in Figure 6-34.

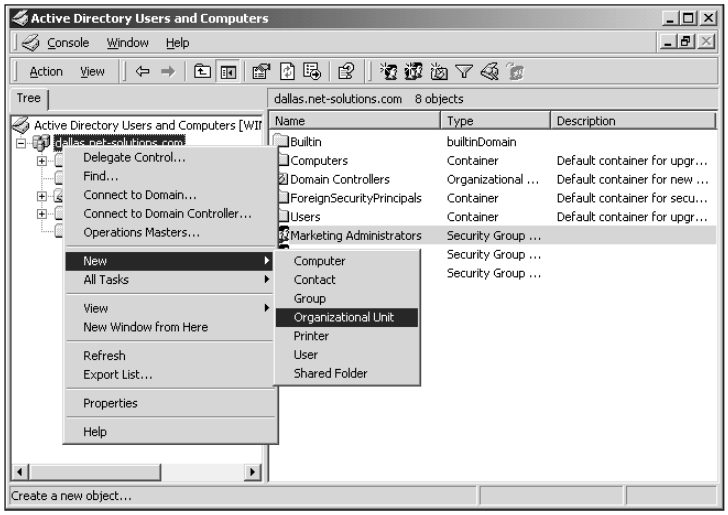

Figure 6-34 Creating a new OU

The first property screen for the new OU asks for a name. This name should be descriptive and should clearly show the role of the OU. Enter it in the Name field, as shown in Figure 6-35.

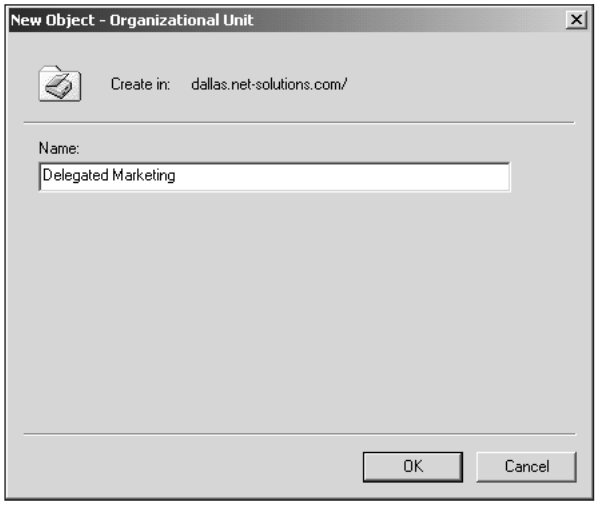

Figure 6-35 Enter the name for the OU

Once the OU is created, it must be populated. To move users, computers, or other objects to an OU, simply open the proper folder and highlight the desired objects. From the context menu, click on Move, as shown in Figure 6-36.

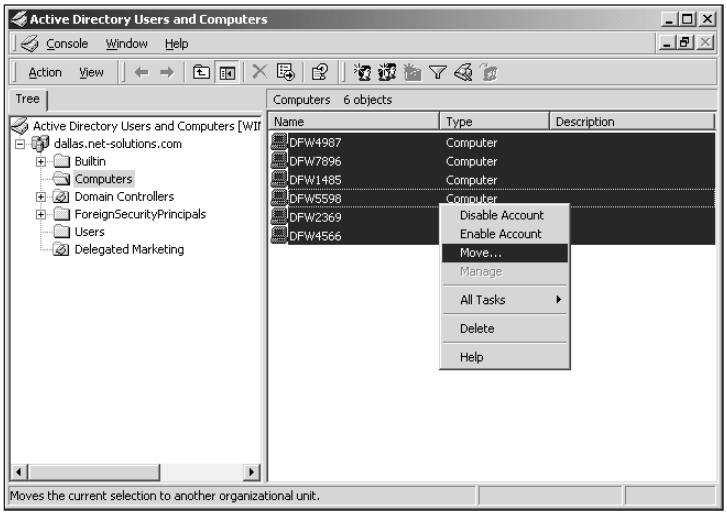

Figure 6-36 Moving objects to an OU

The next step is to select the destination OU for the objects, as shown in Figure 6-37.

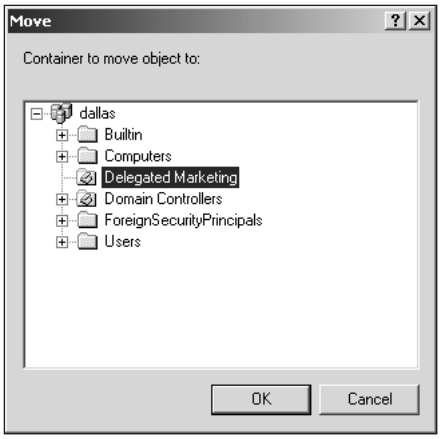

Figure 6-37 Selecting the destination OU

After the various objects are moved into the OU, the contents of that OU can be viewed through the Active Directory Users and Computers console. You can see in Figure 6-38 that we placed in the OU both the Marketing group and the computers the marketing group uses. The computers are placed here for future use; the delegation would work just as well without those objects. The Marketing group contains the user accounts for the marketing department.

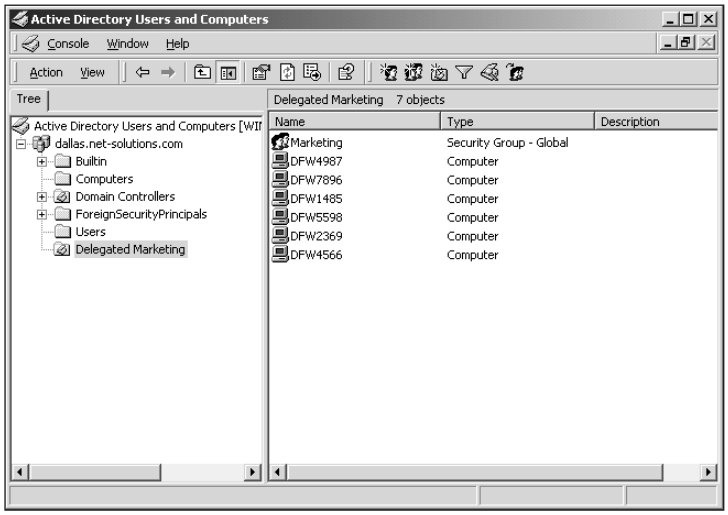

Figure 6-38 Viewing the contents of an OU

Once the OU is created, it is time to delegate control of the OU to a selected few marketing users. Begin by opening the Active Directory Users and Computers console and selecting the desired OU. From the context menu, select Delegate Control, as shown in Figure 6-39.

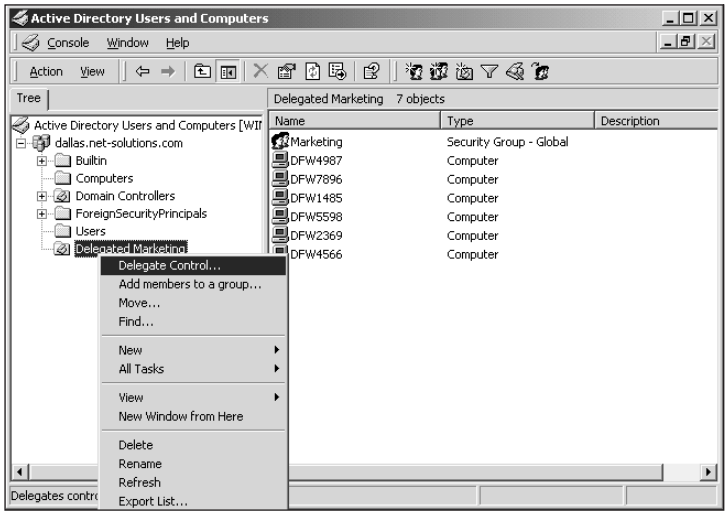

Figure 6-39 Delegating control of an OU

This action will launch the Delegation of Control Wizard, as seen in Figure 6-40. As with most wizards, just click on Next to pass the startup screen.

Figure 6-40 The Delegation of Control Wizard

Next, choose the group and/or users to whom the control is being delegated. In our case, we'll choose a group called Marketing Administrators, as shown in Figure 6-41. This group, which we created earlier, contains the user accounts of the two people trusted to change the passwords.

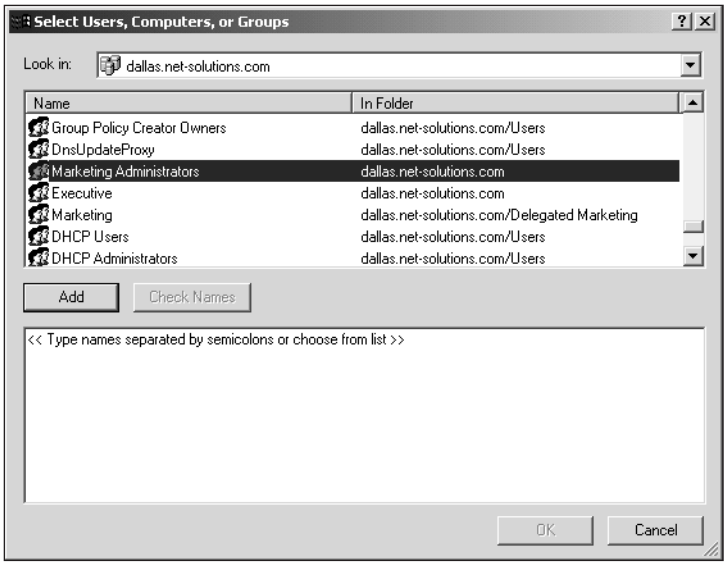

Figure 6-41 Select the group or users to whom control will be delegated

After this selection, choose the rights that the delegate should exercise over the OU. The options you choose here determine the abilities of the delegated administrator. Selecting Reset Passwords On User Accounts will allow the administrators for the OU to reset user passwords. As you can see in Figure 6-42, several other options are available.

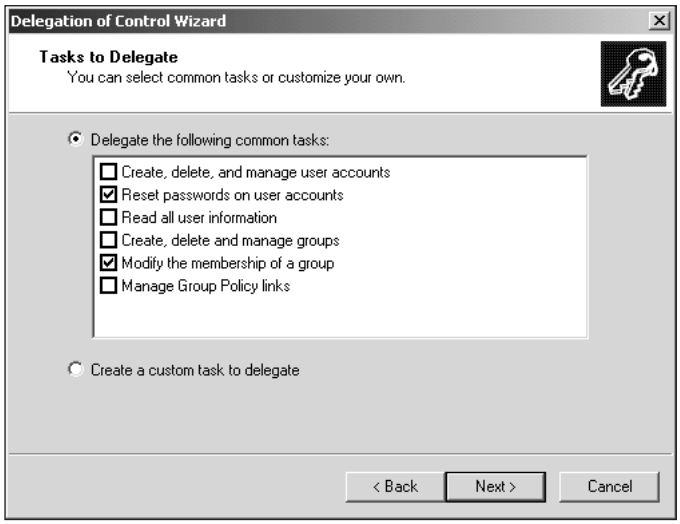

Figure 6-42 Assigning permissions

The last step merely confirms the rights granted to the delegates. You should always double-check and verify that the rights granted actually match the intended purpose. Remember, the rights are inherited throughout the OU. If the rights granted are correct, click on Finish, as shown in Figure 6-43.

Figure 6-43 Verifying the delegated rights

Group Policies and OUs

A second major use of OUs is to assign group policies to particular computers and users. Group policies are used to define default settings for computers and users, such as folder locations, what software can be installed, desktop appearance, and much more. In general, group policies are applied at a domain or site level, but they also can be applied at an OU level in order to generate a specific combination of user and computer environmental factors.

Although the many details of group policies are beyond the scope of this chapter, we will discuss how to associate a new or existing Group Policy with an OU.

First, open the Active Directory Users and Computers management console and navigate to the desired OU. After highlighting the OU, select Properties from the context menu. Doing so will bring up a property sheet like the one shown in Figure 6-44. Select the Group Policy tab to view and modify group policies relating to the OU.

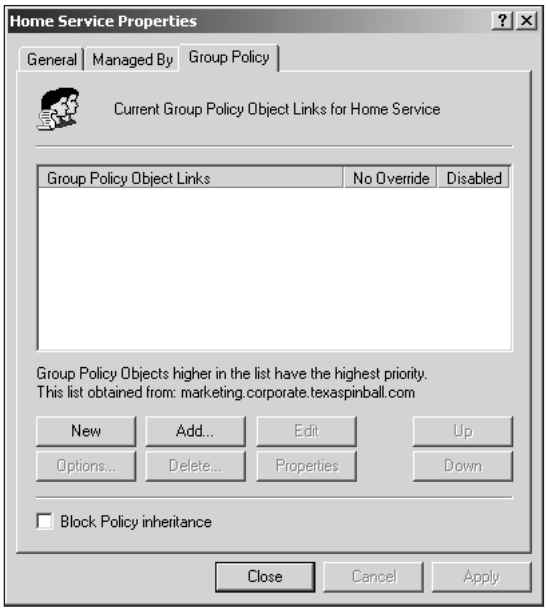

Figure 6-44 Property sheet for Home Service OU

To define a new Group Policy for this OU, begin by clicking on New at the bottom of the screen. A policy called New Group Policy Object will appear. Rename this policy to something memorable and descriptive. In Figure 6-45, the new policy is named Service.

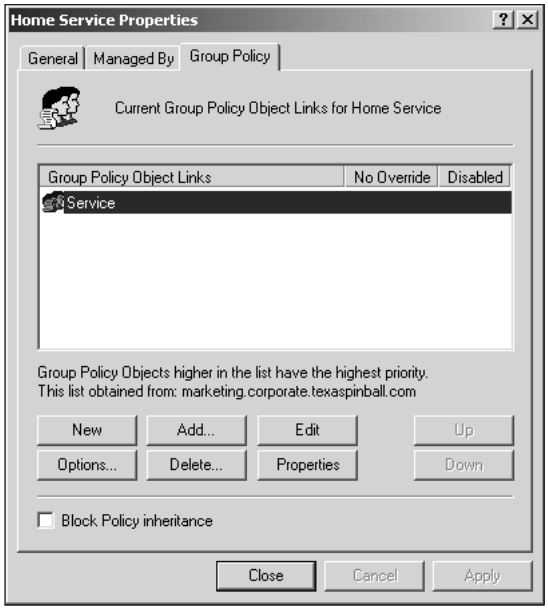

Figure 6-45 The first step in creating a new Group Policy

After the new Group Policy is named, it must be defined. Click on Edit to bring up the Group Policy editor, shown in Figure 6-46. After the policy is modified and saved, clicking on OK on the OU property sheet will finish linking that OU to the policy.

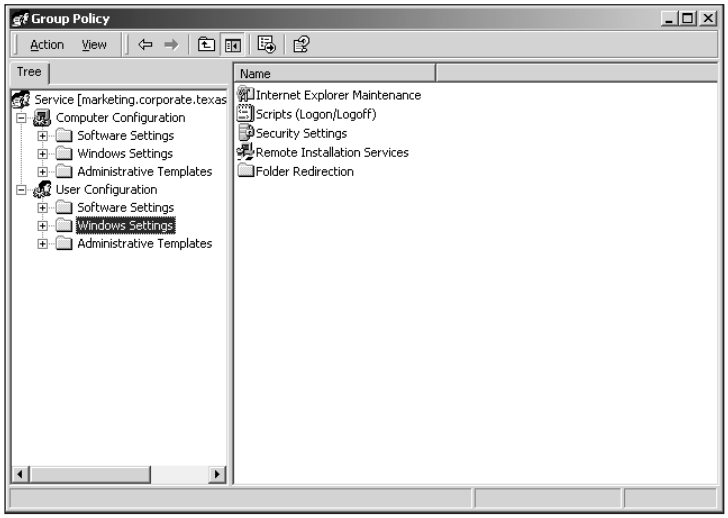

Figure 6-46 The Group Policy editor

OUs can also be linked to group policies that have been previously defined. To link to an existing policy, click on Add from the OU Group Policy properties. Doing so will bring up the screen shown in Figure 6-47. Navigate to the desired policy and select it. Click on OK to link the policy to the OU.

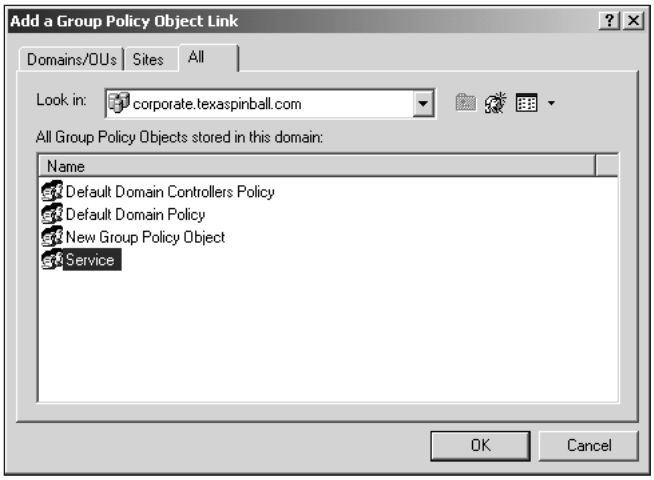

Figure 6-47 Selecting an existing group policy

Naturally, OUs can be unlinked from group policies, as well. To remove a Group Policy from an OU, open the properties of the OU within the Active Directory Users and Computers console. Click on the Group Policy tab, highlight the policy, and click on Delete. Active Directory will offer two options: to unlink the Group Policy from the OU or to delete the Group Policy from the Active Directory environment.

Chapter Summary

- The most important element in designing and implementing a Windows 2000 networking environment remains planning. The Active Directory service requires much more planning than the Windows NT operating system due to elements like sites, site links, and replication scheduling.
- Once the planning is finished, it is time to actually implement the Active Directory structure. The core tool for configuration of the Active Directory services is—not surprisingly—the AD Sites and Services snap-in.
- Sites are areas of high-speed connectivity; typically sites correspond to a physical location and a LAN. Sites are connected via slower links. These links are typically WAN links or, more rarely, heavily utilized LAN links. A site is defined within the AD Sites and Services snap-in through the Sites folder. Windows 2000 creates an initial site called Default-First-Site.

- Once a site is created, an IP subnet or multiple subnets must be associated with the site. You do so through the Subnets folder within the AD Sites and Services snap-in. Now, any new computers created with IP addresses within that subnet are automatically assigned to the site. Subnets are entered with standard dotted octet notation, but are identified in a network/bit-mask format within the Subnets folder.
- Servers created before the sites and subnets are associated are placed in the Default-First-Site. The same is true for servers that have IP addresses outside the ranges associated with particular sites. These servers can be moved between sites through the AD Sites and Services snap-in.
- Sites are linked via site links. A site link is a network connection of some type between at least two sites. As with everything else related to sites, the links are created within the AD Sites and Services snap-in. The cost (relative speed) of the network link can be defined by the network administrator, as well as replication frequency. You also have the option of making the site link active only at certain times. By default, all site links are bridged to allow replication throughout the network environment. If the replication path needs to be controlled, then the default site link bridging can be disabled and specific site link bridges defined. Once specific site link bridges are defined, replication will travel only over the specified links.
- Global Catalog Servers keep a minimal database containing the users and rights from every domain within a forest. The Global Catalog is required to complete the logon process within a native mode, and it also serves to reduce query and logon times. Any DC can be a Global Catalog Server. This arrangement is accomplished by checking the Global Catalog option in the NTDS Settings for that server.
- The multi-master replication process eliminates many of the performance bottlenecks that affect replication in earlier versions of Windows NT, but it leads to additional problems. Active Directory corrects these issues by creating operations masters that have the sole responsibility for several tasks. By default, the first DC created within a forest controls all these roles, but all the functions can be moved to other DCs.
- Organizational Units are groups of users, computers, and other objects that can be administered as a unit. An OU can be used to form a type of informal subdomain within a Windows 2000 domain. You can delegate administrative control of an OU to a particular group or user via the AD Users and Computers console.
- OUs can also be used to assign particular policies to selections of users and computers. You do so through the Group Policy tab in the OU properties. An OU can be linked to a new or existing Group Policy, or unlinked from a policy.